arXiv:1510.03928v2 [q-fin.CP] 26 Mar 2016

Weakly chained matrices, policy iteration, and impulse control

P. Azimzadeh* P. A. Forsyth*

Abstract

This work is motivated by numerical solutions to *Hamilton-Jacobi-Bellman quasi-variational inequalities* (HJBQVIs) associated with *combined stochastic and impulse control* problems. In particular, we consider (i) *direct control*, (ii) *penalized*, and (iii) *semi-Lagrangian* discretization schemes applied to the HJBQVI problem. Scheme (i) takes the form of a *Bellman problem* involving an operator which is not necessarily contractive. We consider the well-posedness of the Bellman problem and give sufficient conditions for convergence of the corresponding policy iteration. To do so, we use *weakly chained diagonally dominant matrices*, which give a graph-theoretic characterization of weakly diagonally dominant M-matrices. We compare schemes (i)–(iii) under the following examples: (a) optimal control of the exchange rate, (b) optimal consumption with fixed and proportional transaction costs, and (c) pricing guaranteed minimum withdrawal benefits in variable annuities. We find that one should abstain from using scheme (i).

Keywords. Hamilton-Jacobi-Bellman equation, combined stochastic and impulse control, policy iteration, weakly chained diagonally dominant matrix, optimal exchange rate, optimal consumption, GMWB

AMS subject classifications. 65N06, 93E20

1 Introduction

This work is motivated by the computation of numerical solutions to *Hamilton-Jacobi-Bellman quasi-variational inequalities* (HJBQVI) associated with combined stochastic and impulse control. These problems are of the form:

Problem. *Find a viscosity solution (see [17, Definition 2.2]) of the HJBQVI*

$$
\begin{aligned}
0 &= F(t,x,u(t,x),Du(t,x),D^2u(t,x)) \\
&:= -\begin{cases} \max\left(\sup_{w\in W}\left\{\dfrac{\partial u}{\partial t}+Ł^w u-\rho u+f^w\right\},\mathcal{M}u-u\right) & \text{on } [0,T)\times(\overline{\Omega}\setminus\Lambda) \\ \max\left(g-u,\mathcal{M}u-u\right) & \text{on } \partial^+\Omega \end{cases}
\end{aligned} \tag{1.1}
$$

*David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1 pazimzad@uwaterloo.ca and paforsyt@uwaterloo.ca

where $\Omega \subset \mathbb{R}^d$ *is open,* $\Lambda \subset \partial\Omega$, $\partial^+\Omega := ([0,T) \times \Lambda) \cup (\{T\} \times \overline{\Omega})$, $Ł^w := Ł(t,x,w)$ *is the (possibly degenerate) generator of an SDE,* $f^w := f(t,x,w)$ *is a forcing term, and* $\mathcal{M}$ *is the impulse (a.k.a. intervention) operator*

$$\mathcal{M}u(t,x) := \sup_{z \in Z(t,x)} \{u(t, x + \Gamma(t,x,z)) + K(t,x,z)\}. \tag{1.2}$$

If $Z(t,x)$ *is empty at a particular point* (t,x), $\mathcal{M}u(t,x)$ *is understood to take the value* $-\infty$, *corresponding to no impulses being allowed at that point.*

We focus on implicit discretization schemes for the HJBQVI problem that do not suffer from the usual timestep restrictions of explicit schemes. In particular, we consider (i) *direct control*, (ii) *penalized* and (iii) *semi-Lagrangian* schemes. The semi-Lagrangian scheme (used for HJBQVIs in [12]) differs from its counterparts in that it handles controlled terms using information from the previous timestep. As such, computing the solution of this scheme does not require an iterative method. However, this scheme requires that the control w in $Ł^w$ appears only in the coefficient of the first-order term. For the other two schemes, an iterative method is needed. The particular iterative method analyzed herein is Howard's *policy iteration* algorithm. Not considered is the alternative *value iteration* algorithm, due to its poor performance as the numerical grid is refined [15, §6.1]. Convergence of policy iteration applied to the penalized scheme turns out to be a trivial consequence of the strict diagonal dominance of the input matrices to policy iteration (§5.2). Convergence of policy iteration applied to the direct control scheme is a more delicate matter, as discussed below.

The direct control scheme takes the form of the fixed point problem

$$\text{find } v \in \mathbb{R}^M \text{ such that } v = \max\left(\sup_{w \in \mathcal{W}} L(w)v + c(w), \sup_{z \in \mathcal{Z}} B(z)v + k(z)\right) \tag{1.3}$$

where $L(w)$ and $B(z)$ are contractive and nonexpansive matrices, respectively. It is understood that the supremum and maximum are element-wise and controls are "row-decoupled" (see §2). [10] gives sufficient conditions for convergence of a policy iteration to the unique solution of (1.3). However, convergence in [10] is conditional on the choice of initial guess [10, Theorem 2 (iii)]. We remove this constraint.

More importantly, [10] restricts the admissible set of controls and imposes a strong assumption on $B(z)$ (of which assumption (H2) in this work is an analogue) to ensure convergence of policy iteration applied to (1.3). Unfortunately, reasonable instances of problem (1.3) (including examples in this work) do not necessarily satisfy this condition. We show that, under a much weaker assumption, a solution to (1.3) is unique. Moreover, when (H2) is not satisfied directly, we provide a way to construct this solution by considering a "modified problem" that satisfies (H2). Roughly speaking, one arrives at the modified problem by removing some suboptimal controls from the control set. However, this procedure is ad hoc (i.e. problem dependent).

To establish the above relaxations, we use *weakly chained diagonally dominant* (WCDD) matrices. WCDD matrices give a graph-theoretic characterization of weakly diagonally dominant M-matrices (Theorem 3.5). The WCDD matrix approach to the convergence of policy iteration applied to (1.3) is intuitive and established using well-known results on policy iteration (Proposition 2.2).

The ad hoc removal of suboptimal controls makes the direct control scheme less robust than its counterparts, for which control sets need not be altered to ensure convergence. It is thus natural to ask if there is an advantage to using a direct control scheme. To answer this, we apply each scheme to the following examples:

- optimal control of the exchange rate;
- optimal consumption with fixed and proportional transaction costs;
- pricing guaranteed minimum withdrawal benefits in variable annuities.

The semi-Lagrangian scheme only requires a single linear solve per timestep since no iterative method is needed. However, as mentioned above, such a scheme cannot be used if the control w appears in the diffusion coefficient of $Ł^w$ (or if the underlying process is Lévy with controlled arrival rate). We find that the penalized scheme performs at least as well the direct control scheme. Both produce nearly identical results and often require roughly the same amount of computation. In the specific case of the optimal consumption problem, the penalized scheme even outperforms the direct control scheme, taking only a few policy iterations to converge per timestep.

We mention that in the infinite-horizon setting ($T = \infty$), optimal consumption with fixed and proportional transaction costs was considered numerically in [11] using iterated optimal stopping, a theoretical tool [24, Chapter 7, Lemma 7.1] for the construction of solutions that has found its way into numerical implementations [20, 5]. Computationally, for finite-horizon problems ($T < \infty$), iterated optimal stopping has high space complexity [3], and is thus not considered here. Also not considered here is the simulation of penalized backward stochastic differential equations [18], a recent alternative well-suited to high-dimensional problems.

In this work, we restrict our attention to problems of dimension three or lower. To keep focus on the interesting aspects of impulse control, we assume that between impulses, the underlying stochastic process associated with the HJBQVI is a Brownian motion with drift $\mu := \mu(t, x, w)$ and scaling $\sigma := \sigma(t, x, w)$ (we can extend to a Lévy process with nontrivial arrival rate by, e.g., [13]). This allows us to write

$$Ł^w u(t, x) := \frac{1}{2} \operatorname{trace}\left(\sigma(t, x, w)\sigma^\intercal(t, x, w) D_x^2 u(t, x)\right) + \langle \mu(t, x, w), D_x u(t, x)\rangle .$$

We mention here that problem (1.3) can also be interpreted as a Bellman problem associated with an infinite-horizon Markov decision process (MDP) with vanishing discount (Example 4.2). In fact, (1.3) is a generalization of a reflecting boundary problem (see, e.g., the monograph of Kushner and Dupuis [19, pg. 39–40]). In the context of MDPs, $L(w)$ and $B(z)$ capture the transition probabilities at states with nonvanishing and vanishing discount factors, respectively. A WCDD matrix condition guarantees the convergence of policy iteration to the unique solution of the Bellman problem (Corollary 4.4). Intuitively, this condition ensures that the underlying Markov chain arrives (with positive probability) at a state with nonvanishing discount independent of the initial state.

We summarize some of our main findings below:

- Policy iteration applied to a (monotone) direct control scheme frequently fails due to the possible singularity of the matrix iterates.
- We establish provably convergent techniques to eliminate singularity. However, applying these techniques is problem-dependent.
- We show that policy iteration applied to a (monotone) penalized scheme never fails. Numerical tests on three problems confirm that such a scheme performs at least as well as (and sometimes better than) its direct control counterpart.

The additional effort required to ensure the convergence in the direct control case along with the comparable (if not better) performance in the penalized case suggests that one should abstain from a direct control scheme.

An outline of this work is as follows. §2 reminds the reader of a well-known result on the convergence of policy iteration. §3 discusses WCDD matrices. §4 gives conditions for the convergence of policy iteration applied to problem (1.3) and its well-posedness under weaker assumptions. A self-contained MDP example is given therein (Example 4.2). §5 introduces numerical schemes for the HJBQVI problem (1.1), with numerical examples given in §6.

2 Policy iteration

In the sequel, we will see that each of the discretization schemes for (1.1) take the form of a Bellman problem:

$$\text{find } v \in \mathbb{R}^M \text{ such that } \sup_{P\in\mathcal{P}} \{-A(P)v + b(P)\} = 0 \tag{2.1}$$

where $A\colon \mathcal{P} \to \mathbb{R}^{M\times M}$ and $b\colon \mathcal{P} \to \mathbb{R}^M$. It is understood that (i) $\mathcal{P} \coloneqq \prod_{i=1}^M \mathcal{P}_i$ is a finite product of nonempty sets, (ii) controls are row-decoupled:

$$[A(P)]_{ij} \text{ and } [b(P)]_i \text{ depend only on } P_i \in \mathcal{P}_i,$$

(iii) the order on $\mathbb{R}^M$ (resp. $\mathbb{R}^{M\times M}$) is element-wise:

$$\text{for } x, y \in \mathbb{R}^M,\ x \geq y \text{ if and only if } x_i \geq y_i \text{ for all } i,$$

and (iv) the supremum is induced by this order:

$$\text{for } \{x(P)\}_{P\in\mathcal{P}} \subset \mathbb{R}^M,\ \sup_{P\in\mathcal{P}} x(P) \text{ is a vector with components } \sup_{P\in\mathcal{P}} [x(P)]_i\,.$$

Let $\textsc{Solve}(A, b, x^0)$ denote a call to a linear solver for $Ax = b$ with initial guess x^0 (algebraically, Solve computes x exactly; in practice, an iterative solver is used and the choice of x^0 affects performance). A *policy iteration* algorithm is given by:

$\textsc{Policy-Iteration}(\mathcal{P}, A(\cdot), b(\cdot), v^0)$

1 **for** $\ell = 1, 2, \ldots$
2 Pick P^ℓ such that $-A(P^\ell)v^{\ell-1} + b(P^\ell) = \sup_{P\in\mathcal{P}}\{-A(P)v^{\ell-1} + b(P)\}$
3 $v^\ell \coloneqq \textsc{Solve}(A(P^\ell), b(P^\ell), v^{\ell-1})$

Definition 2.1 (Monotone matrix)**.** A real square matrix A is monotone (in the sense of Collatz) if for all real vectors v, $Av \geq 0$ implies $v \geq 0$.

We use the following assumptions:

(H0) $P \mapsto A(P)^{-1}$ is bounded on the set $\{P \in \mathcal{P}\colon A(P) \text{ is nonsingular}\}$.

(H1) (i) A and b are bounded and (ii) for all x in $\mathbb{R}^M$, there exists P_x in $\mathcal{P}$ such that $-A(P_x)x + b(P_x) = \sup_{P\in\mathcal{P}}\{-A(P)x + b(P)\}$.

Proposition 2.2 (Convergence of policy iteration)**.** *Suppose (H0), (H1), and that $A(P)$ is a monotone matrix for all P in $\mathcal{P}$. $(v^\ell)_{\ell\geq 1}$ defined by* Policy-Iteration *is nondecreasing and converges to the unique solution v of (2.1). Moreover, if $\mathcal{P}$ is finite, convergence occurs in at most $|\mathcal{P}|$ iterations (i.e. $v^{|\mathcal{P}|} = v^{|\mathcal{P}|+1} = \cdots$).*

The monotone convergence of $(v^\ell)_{\ell\geq 1}$ to the unique solution of (2.1) can be proven similarly to Theorem A.3 of Appendix A. See [7, Theorem 2.1] for a proof of the finite termination when $\mathcal{P}$ is finite.

In practice, $\mathcal{P}$ is often finite, in which case (H0) and (H1) are trivial.

Remark 2.3. Theorem A.3 establishes the existence and uniqueness of solutions to (2.1) independent of (H1.ii). Owing to this, results that rely on Proposition 2.2 can be relaxed to exclude (H1.ii), with the caveat that when $\mathcal{P}$ is infinite, Policy-Iteration be replaced by Sequential-Policy-Iteration (see Appendix A). In this case, the resulting sequence $(v^\ell)_{\ell\geq 1}$ is not necessarily nondecreasing.

3 Weakly chained diagonally dominant matrices

We say row i of a complex matrix $A \coloneqq (a_{ij})$ is strictly diagonally dominant (SDD) if $|a_{ii}| > \sum_{j\neq i} |a_{ij}|$. We say A is SDD if all of its rows are SDD. Weakly diagonally dominant (WDD) is defined with weak inequality instead.

Definition 3.1. A complex square matrix A is said to be a weakly chained diagonally dominant (WCDD) if it satisfies:

(i) A is WDD;

(ii) for each row r, there exists a path in the graph of A from r to an SDD row p.

Recall that the directed graph of an $M \times M$ complex matrix $A \coloneqq (a_{ij})$ is given by the vertices $\{1, \ldots, M\}$ and edges defined as follows: there exists an edge from i to j if and only if $a_{ij} \neq 0$. Note that if r is itself an SDD row, the trivial path $r \to r$ satisfies the requirement of (ii) in the above.

The nonsingularity of WCDD matrices is proven in [27] as a corollary to other results. We provide an elementary proof for the convenience of the reader:

Lemma 3.2. *A WCDD matrix is nonsingular.*

Proof. Suppose $\lambda = 0$ is an eigenvalue of $A \coloneqq (a_{ij})$. Let $v \neq 0$ be an associated eigenvector with components of modulus at most unity. Let r be such that $|v_r| = 1 \geq |v_j|$ for all j. By the Gershgorin circle theorem,

$$|\lambda - a_{rr}| = |a_{rr}| \leq \sum_{j\neq r} |a_{rj}|\,|v_j| \leq \sum_{j\neq r} |a_{rj}|\,.$$

Since A is WDD, it follows that $|a_{rr}| = \sum_{j\neq r} |a_{rj}|$, and hence r is not an SDD row. Therefore, there exists a path $r \to p_1 \to \cdots \to p_k$ where p_k is an SDD row. Since

$$|a_{rr}| = \sum_{j\neq r} |a_{rj}|\,|v_j| = \sum_{j\neq r} |a_{rj}|\,,$$

it follows that $|v_j| = 1$ whenever $|a_{rj}| \neq 0$. Because $|a_{rp_1}| \neq 0$, $|v_{p_1}| = 1$. Repeating the same argument as above with $r = p_1$ yields $|a_{p_1p_1}| = \sum_{j\neq p_1} |a_{p_1j}|$, and hence p_1 is not an SDD row. Continuing the procedure, p_k is not SDD, a contradiction. ∎

We recall some well-known classes of matrices:

Definition 3.3. A Z-matrix is a real matrix with nonpositive off-diagonals.

Definition 3.4. An M-matrix is a monotone Z-matrix.

We are now ready to state a fundamental characterization of WDD M-matrices:

Theorem 3.5 (Characterization theorem)**.** *The following are equivalent:*
(i) A is a WCDD Z-matrix with positive diagonals;
(ii) A is a WDD M-matrix.

Proof. Since a nonsingular WDD Z-matrix with positive diagonals is an M-matrix (a consequence of, e.g., [25, Theorem $1.A_3$]), (i) implies (ii) follows by Lemma 3.2.

As for the converse, since an M-matrix has positive diagonal elements (a consequence of, e.g., [25, Theorem $1.K_{35}$]), it is sufficient to show that any WDD Z-matrix with positive diagonals not satisfying (ii) of Definition 3.1 is singular. Let A be such a matrix. It follows that there exists a WDD row r of A from which there is no path to an SDD row of A. Let $V \coloneqq \{i_1, \dots, i_k\} \subset \{1, \dots, \dim A\}$ be the set of rows reachable from r (note that $r \in V$ since $a_{rr} > 0$). Let v_i be unity if i is in V and zero otherwise. Then, $Av = 0$ and $v \neq 0$, and hence (ii) implies (i). ∎

This characterization is tight: an M-matrix need not be WCDD (e.g. $\left(\begin{smallmatrix} 1 & -2 \\ 0 & 1 \end{smallmatrix}\right)$).

We mention that (i) implies (ii) of Theorem 3.5 appears in [8]. Therein, WCDD Z-matrices with positive diagonals are referred to as *matrices of positive type.* To the authors' best knowledge, the converse does not appear in the literature.

4 The fixed point problem (1.3)

4.1 Convergence of policy iteration

We assume $\mathcal{W} \coloneqq \prod_{i=1}^{M} \mathcal{W}_i$ and $\mathcal{Z} \coloneqq \prod_{i=1}^{M} \mathcal{Z}_i$ appearing in problem (1.3) are finite products of nonempty sets. Let

$$\mathcal{P} \coloneqq \textstyle\prod_{i=1}^{M} \mathcal{P}_i \text{ where } \mathcal{P}_i \coloneqq \mathcal{W}_i \times \mathcal{Z}_i \times \mathcal{D}_i \text{ and } \emptyset \neq \mathcal{D}_i \subset \{0, 1\}\,. \tag{4.1}$$

We associate with each $\psi \coloneqq (\psi_1, \dots, \psi_M)$ in $\mathcal{D} \coloneqq \prod_{i=1}^{M} \mathcal{D}_i$ a diagonal matrix $\Psi \coloneqq \operatorname{diag}(\psi)$. We use ψ and Ψ interchangeably. We write $P \coloneqq (w, z, \psi) \in \mathcal{P}$ where $w \in \mathcal{W}$ and $z \in \mathcal{Z}$. We can transform problem (1.3) into the form (2.1) by taking

$$\begin{aligned} A(P) &\coloneqq (I - \Psi)\,(I - L(w)) + \delta\Psi\,(I - B(z)) \\ b(P) &\coloneqq (I - \Psi)\,c(w) + \delta\Psi k(z) \end{aligned} \tag{4.2}$$

where $\delta = 1$ ($L(w)$ and $B(z)$ are matrices; $c(w)$ and $k(z)$ are vectors). To keep the material general, we henceforth assume the less restrictive condition $\delta > 0$ instead of $\delta = 1$. Before considering the well-posedness of

$$\text{problem (2.1) subject to (4.1) and (4.2)}, \tag{4.3}$$

we establish that the set of solutions to (4.3) is independent of the choice of δ:

Lemma 4.1. *v is a solution of (4.3) with $\delta = 1$ if and only if it is a solution of (4.3) with arbitrary $\delta = \delta_0 > 0$.*

A proof of the above is given in Appendix B. In the sequel, we exploit the fact that policy iteration may converge more rapidly for particular choices of δ. We now visit, as a motivating example, an infinite-horizon MDP with vanishing discount:

Example 4.2. Let $(X^n)_{n\geq 0}$ be a controlled homogeneous Markov chain on a finite state space $\{1, \dots, M\}$. A control at state i is a member of $\mathcal{P}_i$ in (4.1) and written $P_i \coloneqq (w_i, z_i, \psi_i)$. The transition probabilities of the Markov chain are

$$\mathbb{P}(X^{n+1} = j \mid X^n = i) = \begin{cases} w_{ij} & \text{if } \psi_i = 0 \\ z_{ij} & \text{if } \psi_i = 1 \end{cases}$$

where $w_i \coloneqq (w_{i1}, \dots, w_{iM}) \geq 0$ and $w_i e = 1$ (similarly for z_i). That is, members of $\mathcal{W}_i$ and $\mathcal{Z}_i$ are M-dimensional probability vectors. Let

$$v_i \coloneqq \sup_{P\in\mathcal{P}} \mathbb{E}\left[\sum_{n=0}^{\infty} U(X^n, P) \prod_{m=0}^{n-1} D(X^m, P) \,\middle|\, X^0 = i\right] \text{ for all } 1 \leq i \leq M \tag{4.4}$$

where

$$U(i,P) \coloneqq \begin{cases} c_i(w_i) & \text{if } \psi_i = 0 \\ k_i(z_i) & \text{if } \psi_i = 1 \end{cases} \text{ and } D(i,P) \coloneqq \begin{cases} 1/(1+\rho) & \text{if } \psi_i = 0; \\ 1 & \text{if } \psi_i = 1. \end{cases}$$

In the above, $\rho > 0$ is a discount factor. [10, Lemma 5] establishes that the dynamic programming equation associated with (4.4) is exactly (4.3) with $[L(w)]_{ij} \coloneqq w_{ij}/(1+\rho)$, $[B(z)]_{ij} \coloneqq z_{ij}$, $[c(w)]_i \coloneqq c_i(w_i)$, and $[k(z)]_i \coloneqq k_i(z_i)$.

In the above, states i on which $\psi_i = 1$ are the "trouble" states with vanishing discount factor. In fact, requiring $\psi_i = 0$ for all i returns us to a nonvanishing discount factor problem whose well-posedness is easy to establish.

The following assumptions will prove paramount:

(H2) For each $P \coloneqq (w, z, \psi)$ in $\mathcal{P}$ and state i with $\psi_i = 1$, there exists a path in the graph of $B(z)$ from i to some state $j(i)$ with $\psi_{j(i)} = 0$.

(H3) For each $P \coloneqq (w, z, \psi)$ in $\mathcal{P}$, $I - L(w)$ and $I - B(z)$ are Z-matrices with nonnegative diagonals that are SDD and WDD, respectively.

Theorem 4.3. *Suppose (H0)–(H3). $(v^\ell)_{\ell\geq 1}$ defined by* Policy-Iteration *is nondecreasing and converges to the unique solution v of (4.3). Moreover, if $\mathcal{P}$ is finite, convergence occurs in at most $|\mathcal{P}|$ iterations.*

Proof. (H2) and (H3) ensure that $A(P)$ is a WCDD Z-matrix with positive diagonals. The desired result follows from Theorem 3.5 and Proposition 2.2. ∎

Corollary 4.4. *Consider Example 4.2. Suppose (H0)–(H2). $(v^\ell)_{\ell\geq 1}$ defined by* Policy-Iteration *converges to v in $\mathbb{R}^M$ satisfying (4.4).*

An example satisfying (H2) is given:

Figure 4.1.1: Graphs of possible matrices $B(z)$ from Example 4.5

Example 4.5. Consider Example 4.2. Suppose all $P \coloneqq (w, z, \psi)$ in $\mathcal{P}$ satisfy

$$\psi_1 = 0 \text{ and } z_{ij} = 0 \text{ if } 1 < i \leq j. \tag{4.5}$$

This corresponds to (i) a nonvanishing discount at state 1 and that (ii) transitions from a state with vanishing discount are unidirectional (if $\psi_{X^n} = 1$, $X^{n+1} < X^n$ a.s.). See Figure 4.1.1 for example graphs of $B(z)$ subject to (4.5).

4.2 Uniqueness

Let

$$\mathbb{L}v \coloneqq \sup_{w \in \mathcal{W}} \{L(w)v + c(w)\} \text{ and } \mathbb{B}v \coloneqq \sup_{z \in \mathcal{Z}} \{B(z)v + k(z)\}. \tag{4.6}$$

The condition (H2) turns out to be too restrictive for some problems of interest. However, the following weaker property of $\mathbb{B}$ is not unusual:

(H4) For each solution v of (4.3) and each state i, there exist integers $m(i)$ and $n(i)$ such that $0 \leq n(i) < m(i)$ and $[\mathbb{B}^{m(i)}v]_i < [\mathbb{B}^{n(i)}v]_i$.

Lemma 4.6. *Suppose (H4). Let $(P^\ell)_{\ell \geq 0} \coloneqq (w^\ell, z^\ell, \psi^\ell)_{\ell \geq 0}$ be a sequence in $\mathcal{P}$ and v a solution of (4.3) satisfying*

$$-A(P^\ell)v + b(P^\ell) \to \sup_{P \in \mathcal{P}} \{-A(P)v + b(P)\} = 0.$$

There exists $\ell_0 \geq 0$ such that for each $\ell \geq \ell_0$ and state i with $\psi_i^\ell = 1$, there exists a path in the graph of $B(z^\ell)$ from i to some state $j(i, \ell)$ with $\psi_{j(i,\ell)}^\ell = 0$.

Proof. Suppose the contrary. A pigeonhole principle argument yields the existence of a subsequence $(P^{\ell_q})_{q \geq 0} \coloneqq (w^{\ell_q}, z^{\ell_q}, \psi^{\ell_q})_{q \geq 0}$ of $(P^\ell)_{\ell \geq 0}$ such that

- $\psi^{\ell_q} = \psi$ is a constant independent of q;
- the graph of $B(z^{\ell_q})$ (call it G) is a constant independent of q;
- there exists i such that $\psi_i = 1$ and for all j reachable from i (in G), $\psi_j = 1$.

Let $V \coloneqq \{j_1, \ldots, j_k\}$ be the states reachable from i. Let $r \in V \cup \{i\}$ be arbitrary. Since the limit of a convergent sequence equals to the limit of any of its subsequences,

$$\left[B(z^{\ell_q})v + k(z^{\ell_q})\right]_r - v_r = \frac{1}{\delta}\left[-A(P^{\ell_q})v + b(P^{\ell_q})\right]_r \to 0,$$

and hence $[\mathbb{B}v]_r = v_r$. Applying this fact repeatedly,

$$v_i = [\mathbb{B}v]_i = [\mathbb{B}^2 v]_i = \cdots = [\mathbb{B}^{n(i)} v]_i = \cdots = [\mathbb{B}^{m(i)} v]_i.$$

However, $[\mathbb{B}^{m(i)} v]_i < [\mathbb{B}^{n(i)} v]_i$ by (H4), a contradiction. ■

If we take the trivial path $i \to i$ as having length zero, the proof above also implies that for ℓ sufficiently large and for all i, there is a path of length $< m(i)$ (where $m(i)$ is specified by (H4)) in the graph of $B(z^\ell)$ from i to some state $j(i,\ell)$ with $\psi^\ell_{j(i,\ell)} = 0$. An example is given below:

Example 4.7. Consider Example 4.2 with $\mathcal{Z}_i = \mathcal{Z}_j$ for all states i and j. For all states i, let $k_i(z_i) \coloneqq -C < 0$. It follows that for all x in $\mathbb{R}^M$,

$$\mathbb{B}^2 x = \sup_{z,z' \in \mathcal{Z}} \{B(z)B(z')x\} - 2C < \sup_{z \in \mathcal{Z}} \{B(z)x\} - C = \mathbb{B}x,$$

so that (H4) is satisfied with $1 = n(i) < m(i) = 2$ for all i. Intuitively, the controller pays twice the cost to apply $\mathbb{B}$ twice.

In this case, denoting by v a solution of (4.3), the control shown in Figure 4.1.1a cannot correspond to some P_v satisfying $-A(P_v)v + b(P_v) = \sup_{P \in \mathcal{P}}\{-A(P)v + b(P)\} = 0$ since any path from $i > 2$ to $j = 1$ is of length at least $m(i) = 2$.

We can now prove uniqueness independent of (H2):

Theorem 4.8. *Suppose (H0), (H3), and (H4). A solution of (4.3) is unique.*

Proof. Let x and y be two solutions and $(P^\ell)_{\ell \geq 0}$ be a sequence in $\mathcal{P}$ such that

$$-A(P^\ell)y + b(P^\ell) \to \sup_{P \in \mathcal{P}} \{-A(P)y + b(P)\} = 0.$$

It follows from (H3), (H4), and Lemma 4.6 that we can, without loss of generality, assume $A(P^\ell)$ is a WCDD Z-matrix with positive diagonals, and hence an M-matrix by Theorem 3.5. For some sequence $(\epsilon^\ell)_{\ell \geq 0}$ in $\mathbb{R}^M$ with $\epsilon^\ell \to 0$, we can write

$$-A(P^\ell)y + b(P^\ell) + \epsilon^\ell = 0 = \sup_{P \in \mathcal{P}} \{-A(P)x + b(P)\} \geq -A(P^\ell)x + b(P^\ell),$$

so that $A(P^\ell)(x - y) \geq -\epsilon^\ell$. Since the inverse of a monotone matrix has nonnegative elements and $P \mapsto A(P)^{-1}$ is bounded by (H0), $x - y \geq 0$. Similarly, $y - x \geq 0$. ■

Unfortunately, the conditions of Theorem 4.8 cannot guarantee that the iterates $(v^\ell)_{\ell \geq 1}$ given by policy iteration are well-defined, as $A(P^\ell)$ may be singular for some $\ell \geq 1$. This is demonstrated in the following example, for which (H2) does not hold:

Example 4.9 (Failure of policy iteration)**.** Consider Example 4.2. For all states i, let $\mathcal{Z}_i \coloneqq \{e^j\}_{j=1}^M$ be the set of standard basis vectors and

$$k_i(z_i) \coloneqq -C - \sum_j z_{ij}\,|i - j| < 0 \text{ where } C > 0.$$

As in Example 4.7, (H4) is satisfied due to the fixed cost C.

Let $\delta := 1$. Suppose there exists a state r with $1 \in \mathcal{D}_r$ and $c_r(w_r) < -C$ for all controls w in $\mathcal{W}$. It is readily verified that POLICY-ITERATION initialized with the zero vector $v^0 := 0$ picks $P^1 := (w^1, z^1, \psi^1)$ with $z^1_r = e^r$ and $\psi^1_r = 1$. It follows that

$$[A(P^1)]_{rj} = [I - B(z^1)]_{rj} = [I]_{rj} - [I]_{rj} = 0 \text{ for all } j$$

so that $A(P^1)$ is singular, and hence v^1 is undefined.

For any $\ell \geq 1$, it is possible to construct more complicated examples in which the matrices $A(P^1), \ldots, A(P^{\ell-1})$ are nonsingular while $A(P^\ell)$ is singular. That is, policy iteration can fail at any iterate.

4.3 Policy iteration on a modified problem

As demonstrated in the previous section, if (H2) is not satisfied, policy iteration may fail. We may, however, hope to construct a solution by performing policy iteration on a "modified problem" with control set $\mathcal{P}'$ obtained by removing controls P in $\mathcal{P}$ that render $A(P)$ singular.

We define (H1)′ by replacing all occurrences of $\mathcal{P}$ with $\mathcal{P}'$ in the definition of (H1). (H2)′ and (H3)′ are defined similarly. We can now state the above idea precisely:

Theorem 4.10. *Let $\mathcal{P}' := \prod_{i=1}^M \mathcal{P}'_i$ where each $\mathcal{P}'_i \subset \mathcal{P}_i$ is nonempty. Suppose (H0), (H1)′, (H2)′, (H3), (H4), and*

$$\text{for all } v \text{ in } \mathbb{R}^M, \quad \sup_{P \in \mathcal{P}'} \{-A(P)v + b(P)\} = 0 \implies \sup_{P \in \mathcal{P}} \{-A(P)v + b(P)\} = 0. \tag{4.7}$$

$(v^\ell)_{\ell \geq 1}$ defined by POLICY-ITERATION$(\mathcal{P}', A(\cdot), b(\cdot), v^0)$ *is nondecreasing and converges to the unique solution of (4.3). Moreover, if $\mathcal{P}'$ is finite, convergence occurs in at most $|\mathcal{P}'|$ iterations.*

Proof. Since $\mathcal{P}' \subset \mathcal{P}$, it follows immediately that (H0)′ and (H3)′ are satisfied, so that by Theorem 4.3, $(v^\ell)_{\ell \geq 1}$ is well-defined and converges to the unique solution v of the modified problem. That is, $v^\ell \to v$ and $\sup_{P \in \mathcal{P}'} \{-A(P)v + b(P)\} = 0$. By (4.7), v solves the original problem (4.3). Since solutions to (4.3) are unique by Theorem 4.8, the desired result follows. ∎

We now give a nontrivial example (in the sense that (H2) fails) for which we can apply Theorem 4.10:

Example 4.11. Consider Example 4.2. Let (i) $\mathcal{Z}$ and k be given as in Example 4.9, (ii) $\mathcal{W}_i$ be the set of all M-dimensional probability vectors w_i with $w_{ij} = 0$ whenever $|i - j| > 1$, (iii) c be continuous and bounded, and (iv) $\overline{c} := \max_{w \in \mathcal{W}} c(w)$ with $\overline{c}_{i-1} \geq \overline{c}_i$ for all $1 < i \leq M$. Let $\mathcal{P}'$ be all $P := (w, z, \psi)$ in $\mathcal{P}$ satisfying (4.5). Then, the conditions of Theorem 4.10 are satisfied.

Proof. It is straightforward to verify (H0), (H1)′, (H2)′, (H3), and (H4). Thus, it is sufficient to show (4.7). We write $\mathcal{P}'_i := \mathcal{W}_i \times \mathcal{Z}'_i \times \mathcal{D}'_i$ and define $[\mathbb{B}'x]_i$ for $i > 1$ by replacing $\mathcal{Z}$ with $\mathcal{Z}' := \prod_{i=1}^M \mathcal{Z}'_i$ in (4.6).

We first show that the solution v to the modified problem is nonincreasing:

$$v_{i-1} \geq v_i \text{ for all } 1 < i \leq M.$$

Suppose the contrary. Let $r > 1$ be the minimal element such that $v_{r-1} < v_r$. If $v_r = [\mathbb{B}'v]_r$, then $v_r = v_j - C - |r-j|$ for some $j < r$. Either $j = r-1$ or

$$v_{r-1} \geq [\mathbb{B}'v]_{r-1} \geq v_j - C - |(r-1)-j| \geq v_r$$

(both are contradictions). It follows that $v_r = [\mathbb{L}v]_r$. Letting $w_{0j}, v_0 \coloneqq 0$ for notational convenience, assumption (ii) implies

$$v_{r-1} \geq [\mathbb{L}v]_{r-1} = \max_{w \in \mathcal{W}} \left\{ \sum_{j=r-2}^{r} \frac{w_{r-1,j}}{1+\rho} v_j + [c(w)]_{r-1} \right\} \geq \frac{v_{r-1}}{1+\rho} + \bar{c}_{r-1}$$

so that $v_{r-1} \geq (1+1/\rho)\bar{c}_{r-1}$. If $v_r \geq v_{r+1}$, it follows similarly from $v_r = [\mathbb{L}v]_r$ that $v_r \leq (1+1/\rho)\bar{c}_r$ so that $v_{r-1} \geq v_r$ (a contradiction) and hence it must be the case that $v_r < v_{r+1}$. We can repeat this argument inductively to arrive at the contradiction

$$v_{r-1} < v_r < \cdots < v_M \text{ and } v_{r-1} \geq (1+1/\rho)\bar{c}_{r-1} \geq (1+1/\rho)\bar{c}_M \geq v_M.$$

Since v is nonincreasing, $v \geq \mathbb{B}v$ (it is suboptimal to take $\psi_i = 1$ and $z_{ij} = 1$ for states i and j with $j \geq i$), and hence (4.7) holds. ■

5 Numerical schemes for the HJBQVI problem

All numerical schemes herein are on a rectilinear grid

$$\{t^1, \ldots, t^N\} \times \{x_1^1, x_2^1, \ldots\} \times \cdots \times \{x_1^d, x_2^d, \ldots\}$$

where $0 =: t^1 < \cdots < t^N \coloneqq T$ and $x_1^j < x_2^j < \cdots$ for all j. Multi-indices are used (i.e. $x_i \coloneqq (x_{i_1}, \ldots, x_{i_d})$). M denotes the number of spatial points x_i. For functions $q \coloneqq q(t,x)$ defined on $[0,T] \times \mathbb{R}^d$, the shorthands $q_i^n \coloneqq q(t^n, x_i)$ and $q^n(x) \coloneqq q(t^n, x)$ are employed. In the absence of ambiguity, we use q^n to denote the vector with components q_i^n and take $\Delta t \coloneqq t^{n+1} - t^n$. It is understood that $\max_n\{t^{n+1} - t^n\} \to 0$ and $\max_i\{x_{i+1} - x_i\} \to 0$ as $h \to 0$, where h denotes a "global" discretization parameter that controls the coarseness of the grid.

Control sets W and $Z(t,x)$ are approximated by finite sets $\emptyset \neq W_h \subset W$ and $Z_h(t,x) \subset Z(t,x)$. The reader concerned with consistency should impose some regularity to justify this approximation, such as: (i) W is compact, (ii) Z is everywhere compact and continuous with respect to the Hausdorff metric, and (iii) $\max_{w \in W} \min_{w_h \in W_h} |w - w_h| \to 0$ as the discretization parameter $h \to 0$ along with an identical pointwise condition for Z and Z_h.

The discretized impulse operator (1.2) is

$$[\mathcal{M}_h^n u^n]_i \coloneqq \max_{z \in (Z_h)_i^n} \{u^n[\![x_i + \Gamma_i^n(z)]\!] + K_i^n(z)\}$$

where $\varphi[\![x]\!]$ denotes linear interpolation using the value of φ on grid nodes. It is understood that controls z that cause $x_i + \Gamma_i^n(z)$ to exit the numerical grid are not included in $(Z_h)_i^n$. We use $Ł_h^n(w)$ to denote a consistent discretization of $Ł^w$ with coefficients frozen at $t = t^n$.

Recall that in (1.1), $\Lambda \subset \partial\Omega$ is a special subset of the boundary at which a Dirichlet-like condition is applied. To distinguish points, we denote by Φ a diagonal matrix satisfying $[\Phi]_{ii} = 0$ whenever x_i is in Λ and $[\Phi]_{ii} = 1$ otherwise.

Since the Dirichlet-like condition is imposed at the final time $t = T$, the numerical method proceeds backwards in time (i.e. from t^{n+1} to t^n). More precisely, letting $u_i^N \coloneqq g_i^N$, the numerical solution u^n at timestep $1 \le n < N$ produced by each scheme (given the solution at the previous timestep, u^{n+1}) is written as a solution of (2.1) with A and b picked appropriately. Control sets are given by (4.1) and

$$\mathcal{W}_i \coloneqq W_h,\ \mathcal{Z}_i \coloneqq \begin{cases} (Z_h)_i^{n_0} & \text{if } (Z_h)_i^{n_0} \neq \emptyset \\ \{\emptyset\} & \text{otherwise} \end{cases}, \text{ and } \mathcal{D}_i \coloneqq \begin{cases} \{0,1\} & \text{if } (Z_h)_i^{n_0} \neq \emptyset \\ \{0\} & \text{otherwise} \end{cases} \tag{5.1}$$

where n_0 is $n+1$ for the semi-Lagrangian scheme (see §5.3) and n otherwise. As a technical detail, we take $\mathcal{Z}_i$ to be a nonempty set (we choose $\{\emptyset\}$ arbitrarily) whenever $(Z_h)_i^{n_0}$ is empty to ensure that the product $\mathcal{W}_i \times \mathcal{Z}_i \times \mathcal{D}_i$ of (4.1) is nonempty.

We make the following assumptions:

(A0) $\mathcal{W} \coloneqq \prod_{i=1}^M \mathcal{W}_i$ and $\mathcal{Z} \coloneqq \prod_{i=1}^M \mathcal{Z}_i$ are finite.

(A1) For all w in $\mathcal{W}$, $-Ł_h^n(w)$ is a WDD Z-matrix with nonnegative diagonals.

(A2) For all z in $\mathcal{Z}$, $B^n(z)$ is a right Markov matrix with $[B^n(z)v]_i = v[\![x_i + \Gamma_i^n(z_i)]\!]$.

(A3) $\rho \ge 0$ and $\delta, \epsilon > 0$.

Since (A0) ensures that $\mathcal{P}$ is finite, all schemes in the sequel satisfy (H0) and (H1).

Remark 5.1. Barles and Souganidis [4] prove that a numerical scheme converges to the unique viscosity solution of a fully nonlinear second order equation (such as (1.1)) satisfying a comparison result if it is monotone in the viscosity sense, ℓ_∞ stable, and consistent. Comparison results for the HJBQVI (1.1) are provided in [26, Theorem 5.11]. (A1) and (A2) ensure monotonicity (see [23, Section 1.3] for an example of a stable nonmonotone scheme that fails to converge). For brevity, we do not give proofs of consistency or discuss stability here.

5.1 Direct control

In a direct control formulation, either the generator ($\sup_{w \in W}\{\partial u/\partial t + Ł^w u - \rho u + f^w\}$) or impulse ($\mathcal{M}u - u$) component is active at any grid point. Since these have different units, comparing them in floating point arithmetic requires a scaling factor $\delta > 0$ to ensure fast convergence [16] (see also Lemma 4.1). Scaling by δ and discretizing (1.1) (ignoring boundary conditions) yields

$$\max\left(\max_{w \in W_h}\left\{\frac{u_i^{n+1} - u_i^n}{\Delta t} + [Ł_h^n(w)u^n]_i - \rho u_i^n + f_i^n(w)\right\}, \delta\left([\mathcal{M}_h^n u^n]_i - u_i^n\right)\right) = 0.$$

Including boundary conditions, this is put in the form of (4.3) by taking

$$\begin{aligned} L(w) &\coloneqq \Phi\left(Ł_h^n(w) - \rho I\right)\Delta t; & c(w) &\coloneqq \Phi\left(u^{n+1} + f^n(w)\Delta t\right) + (I - \Phi)\, g^n; \\ B(z) &\coloneqq B^n(z); & k(z) &\coloneqq K^n(z). \end{aligned} \tag{5.2}$$

With B and k given above, the operator $\mathcal{M}_h^n$ is equivalent to $\mathbb{B}$ defined in (4.6).

L and B given above satisfy (H3) due to (A1)–(A3). Therefore, (H4) is a sufficient condition for uniqueness of solutions (Theorem 4.8). Similarly, (H2) is a sufficient condition for convergence of the corresponding policy iteration (Theorem 4.3).

5.2 Penalized

A penalized formulation (treated in detail in [28]) imposes a penalty scaled by $1/\epsilon' \gg 0$ whenever $\mathcal{M}u > u$. The scheme is given by:

$$\max_{w \in W_h} \left\{ \frac{u_i^{n+1} - u_i^n}{\Delta t} + \left[Ł_h^n(w) u^n \right]_i - \rho u_i^n + f_i^n(w) \right\} + \max\left(\left[\mathcal{M}_h^n u^n \right]_i - u_i^n, 0 \right) / \epsilon' = 0.$$

For simplicity, we take $\epsilon' \coloneqq \epsilon \Delta t$ for some $\epsilon > 0$. Including boundary conditions, this is put in the form (2.1) by taking

$$\begin{aligned} A(P) &\coloneqq I + \Phi \left(\rho I - Ł_h^n(w) \right) \Delta t + \Psi \left(I - B^n(z) \right) / \epsilon; \\ b(P) &\coloneqq \Phi \left(u^{n+1} + f^n(w) \Delta t \right) + \left(I - \Phi \right) g^n + \Psi K^n(z) / \epsilon. \end{aligned}$$

Convergence of the corresponding policy iteration is trivial since $A(P)$ is an SDD Z-matrix with positive diagonals (by virtue of (A1)–(A3)), and hence an M-matrix.

5.3 Semi-Lagrangian

The crux of a semi-Lagrangian scheme is the use of a *Lagrangian derivative* to remove the D_x coefficient's dependency on the control w. It is assumed that (i) σ is independent of the control and (ii) the drift μ and forcing term f can be split into (sufficiently regular) controlled and uncontrolled components:

$$\mu(y, w) = \hat{\mu}(y) + \hat{\hat{\mu}}(y, w) \text{ and } f(y, w) = \hat{f}(y) + \hat{\hat{f}}(y, w).$$

We now give some intuition behind a semi-Lagrangian scheme. Consider a generator $\hat{Ł}$ corresponding to an uncontrolled SDE:

$$\hat{Ł} u(y) \coloneqq Ł(w) u(y) - \left\langle \hat{\hat{\mu}}(y, w), D_x u(y) \right\rangle.$$

Letting $X \coloneqq X(t)$ denote a d-dimensional trajectory satisfying

$$X(t^n) = x_i \text{ and } dX(t) = \hat{\hat{\mu}}(t, X(t), w) dt \text{ on } (t^n, t^{n+1}]$$

so that $X(t^{n+1}) \approx X(t^n) + \hat{\hat{\mu}}(t^n, X(t^n), w) \Delta t = x_i + \hat{\hat{\mu}}_i^n(w) \Delta t$, we define the Lagrangian derivative with respect to X as

$$\frac{Du}{Dt}(t, X(t), w) \coloneqq \frac{\partial}{\partial t} \left[u(t, X(t)) \right] = \frac{\partial u}{\partial t}(t, X(t)) + \left\langle \hat{\hat{\mu}}(t, X(t), w), D_x u(t, X(t)) \right\rangle.$$

Ignoring boundary conditions, we substitute $\dfrac{Du}{Dt}$ into (1.1) to get

$$\max\left(\sup_{w \in W} \left\{ \frac{Du}{Dt}^w + \hat{Ł} u - \rho u + f^w \right\}, \mathcal{M} u - u \right) = 0.$$

A discretization of the above is

$$\begin{aligned} \max\Bigg(\max_{w \in W_h} \left\{ u^{n+1} \left[\!\!\left[x_i + \hat{\hat{\mu}}_i^n(w) \Delta t \right]\!\!\right] + \hat{\hat{f}}_i^{n+1}(w) \Delta t \right\}, \left[\mathcal{M}_h^{n+1} u^{n+1} \right]_i \Bigg) \\ - u_i^n + \left(\left[\hat{Ł}_h^n u^n \right]_i - \rho u_i^n + \hat{f}_i^n \right) \Delta t = 0. \end{aligned}$$

It is understood that controls w that cause $x_i + \hat{\mu}_i^n(w)\Delta t$ to exit the numerical grid are not considered at node i. Consistency of this scheme (subject to some mild assumptions) can be shown similarly to [12, Lemma 6.6].

In lieu of (A1), we assume:

(A1′) $-\hat{\mathrm{L}}_h^n$ is a WDD Z-matrix with nonnegative diagonals.

Let $\vec{x}$ denote a vector with components x_i. Including boundary conditions, this is put in the form (2.1) by taking

$$\begin{aligned} A &\coloneqq I + \Phi\left(\rho I - \hat{\mathrm{L}}_h^n\right)\Delta t; \\ b(P) &\coloneqq \Phi\left(\hat{f}^n\Delta t + (I-\Psi)\left(u^{n+1}\left[\!\left[\vec{x} + \hat{\mu}^n(w)\Delta t\right]\!\right] + \hat{\hat{f}}^{n+1}(w)\Delta t\right)\right) \\ &\qquad + (I-\Phi)\, g^n + \Psi\left(B^{n+1}(z)u^{n+1} + K^{n+1}(z)\right). \end{aligned}$$

Since A is independent of P, (2.1) becomes $Av = \max_{P\in\mathcal{P}}\{b(P)\}$; no iterative method is required. A is nonsingular since it is SDD (by virtue of (A1′) and (A3)).

6 Examples

The remainder of this work focuses on numerical examples.

6.1 Optimal combined control of the exchange rate

The following is studied in [22, 9]. Consider a government able to influence the foreign exchange (FEX) rate of its currency by:

- choosing the domestic interest rate **(stochastic control)**;
- buying or selling foreign currency **(impulse control)**.

Let $(r_t)_{t\geq 0}$ denote the domestic interest rate process and $\overline{r}$ the foreign interest rate. At any point in time, the government can buy ($z > 0$) or sell ($z < 0$) foreign currency to influence the FEX market. $(X_t)_{t\geq 0}$, the log of the FEX rate, follows

$$\begin{aligned} dX_t &= -a\,(r_t - \overline{r})\,dt + \sigma d\mathfrak{W}_t \qquad \text{if } \tau_j < t < \tau_{j+1} \qquad \text{(stochastic control)}; \\ X_{\tau_{j+1}} &= X_{\tau_{j+1}-} + z_{\tau_{j+1}} \qquad\qquad\qquad\qquad\qquad\quad \text{(impulse control)}. \end{aligned}$$

$(\mathfrak{W}_t)_{t\geq 0}$ is a standard Brownian motion. $a > 0$ parameterizes the effect of the interest rate differential, $w_t \coloneqq r_t - \overline{r}$, on the FEX rate.

Let $\theta \coloneqq (w, \tau_1, \tau_2, \ldots, z_{\tau_1}, z_{\tau_2}, \ldots)$ where (i) $(w_t)_{t\geq 0}$ is an adapted process, (ii) $\tau_1, \tau_2, \ldots$ are stopping times with $0 =: \tau_0 \leq \tau_1 \leq \tau_2 \leq \ldots \leq T$, and (iii) z_{τ_k} is a τ_k-measurable random variable taking values from some set $Z(\tau_k, X_{\tau_k})$. Any such θ satisfying these properties is referred to as a *combined control*.

A combined control is *admissible* if at all times t, $w_{\min} \leq w_t \leq w_{\max}$ (alternatively, we could impose this up to null sets). Let Θ denote the set of all admissible controls. The optimal cost at time t when $X_t = x$ is given by

$$u(t,x) \coloneqq e^{\rho t}\sup_{\theta\in\Theta}\mathbb{E}^{(t,x)}\left[-\int_t^T e^{-\rho s}\left(p(X_s) + bw_s^2\right)ds - \sum_{\tau_j\leq T} e^{-\rho\tau_j}\left(\lambda\left|z_{\tau_j}\right| + C\right)\right]. \tag{6.1}$$

Parameter		Value	
Discount factor	ρ	2%	per annum
Volatility	σ	30%	per annum
Expiry	T	10	years
Optimal parity	$x^\star$	0	
Interest rate differential	W	0-7%	per annum
Interest rate differential effect	a	0.25	
Interest rate differential cost	b	3	
Scaled transaction cost	λ	1	
Fixed transaction cost	C	0.1	

Table 6.1.1: Optimal combined control of the exchange rate: parameters

The cost of the distance of the FEX rate to the optimal parity $x^\star$ is parameterized by the function p. We take $p(x) \coloneqq (\max(x - x^\star, 0))^2$. The constant $b \geq 0$ parameterizes the cost associated with a nonzero interest rate differential. $\lambda \geq 0$ and $C > 0$ parameterize the cost of an impulse. $\rho \geq 0$ is a discount factor.

It is well-known [6] that the dynamic programming equation associated to (6.1) is the HJBQVI on $\Omega \coloneqq \mathbb{R}$ and $\Lambda \coloneqq \emptyset$ given by (1.1) with $g(T, x) \coloneqq 0$ and

$$\begin{aligned} W &\coloneqq [w_{\min}, w_{\max}]; & Z(t,x) &\coloneqq \mathbb{R}; \\ Ł(t,x,w) &\coloneqq -aw\frac{\partial u}{\partial x} + \frac{1}{2}\sigma^2\frac{\partial^2 u}{\partial x^2}; & \Gamma(t,x,z) &\coloneqq z; \\ f(t,x,w) &\coloneqq -p(x) - bw^2; & K(t,x,z) &\coloneqq -\lambda|z| - C. \end{aligned}$$

6.1.1 Convergence of the direct control scheme

Discretization requires that we truncate $[0,T] \times \mathbb{R}$ to $[0,T] \times [x_1, x_M]$ and $Z(t,x) = \mathbb{R}$ to $[x_1, x_M] - x$ so that the exchange rate after an impulse, $x + \Gamma(t,x,z) = x + z$, remains in the computational domain. Let $\Delta z > 0$ divide $x_M - x_1$. A discretization of the truncated $Z(t,x)$ is

$$(Z_h^n)_i \coloneqq \{0, \Delta z, 2\Delta z, \ldots, x_M - x_1\} + (x_1 - x_i)\,.$$

An artificial Neumann boundary condition $\partial u/\partial x = 0$ is used at x_1 and x_M so that the first and last rows of $Ł_h^n(w)$ are zero. In particular, we assume an upwind three-point stencil [15, Appendix C] so that

$$[Ł_h^n(w)v]_i \coloneqq \begin{cases} 0 & \text{if } i = 1 \text{ or } i = M \\ (v_{i-1} - v_i)\,\alpha_i^n(w) + (v_{i+1} - v_i)\,\beta_i^n(w) & \text{otherwise} \end{cases}$$

where $\alpha_i^n(w)$ and $\beta_i^n(w)$ are nonnegative constants arising from the discretization.

The direct control problem is given by (4.3) subject to (5.1) and (5.2). It is easy to verify that $\mathbb{B}^2 x < \mathbb{B}x$ for all x so that (H4) is satisfied (recall $\mathbb{B} = \mathcal{M}_h^n$). By Theorem 4.8, solutions to the problem are unique. However, policy iteration may fail since (H2) is not satisfied. A trivial example violating (H2) is that of a cycle between two nodes $x_i \neq x_j$ (e.g. $x_i + \Gamma(t, x_i, x_j - x_i) = x_j$ and $x_j + \Gamma(t, x_j, x_i - x_j) = x_i$).

We perform policy iteration on a modified problem with control set $\mathcal{P}' \subsetneq \mathcal{P}$ consisting of all controls $P \coloneqq (w, z, \psi)$ in $\mathcal{P}$ satisfying

$$\psi_1 = 0 \text{ and } z_i < 0 \text{ for all } i > 1$$

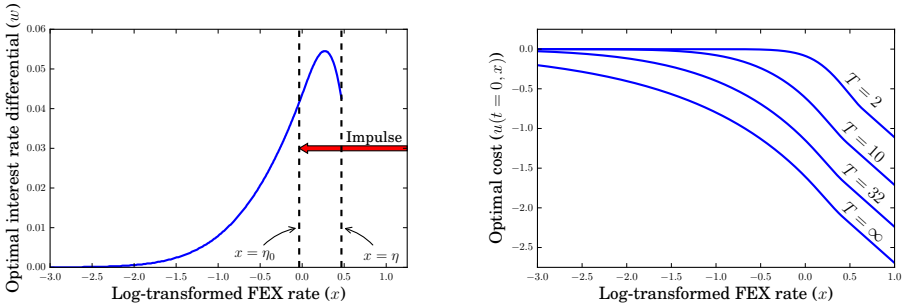

(a) Optimal control ($T = 10$)

(b) Values for several times-to-expiry

Figure 6.1.1: Optimal combined control of the exchange rate at initial time

h	x nodes	w nodes	z nodes	Timesteps
1	32	8	16	16
1/2	64	16	32	32
⋮	⋮	⋮	⋮	⋮

Table 6.1.2: Optimal combined control of the exchange rate: numerical grid

so that (H2)′ holds. If u^{n+1} is nonincreasing (i.e. $u_{i-1}^{n+1} \geq u_i^{n+1}$), we can use the same arguments as in Example 4.11 to establish that the solution $v = u^n$ of the modified problem solves the original problem (i.e. (4.7) is satisfied) and is nonincreasing. Since $u^N = 0$ is nonincreasing, induction yields convergence of the scheme at each timestep.

Remark 6.1. The condition $z_i < 0$ appeals to intuition: the domestic government should never perform an impulse that weakens the domestic currency (i.e. $z_i \geq 0$).

6.1.2 Optimal control

If the currency is sufficiently weak, the government intervenes in the FEX market. That is, at time t, the impulse occurs only on $[\eta(t), \infty)$ for some $\eta(t)$ (the region $(-\infty, \eta(t))$ on which the impulse is not applied is referred to as the *continuation region*, corresponding to nodes i with $\psi_i = 0$ in the numerical solution). When the FEX rate at time t enters $[\eta(t), \infty)$, the government intervenes to bring it back to $\eta_0(t) < \eta(t)$. This phenomenon is shown in Figure 6.1.1a. The optimal cost u for varying expiry times T is shown in Figure 6.1.1b.

6.1.3 Convergence tests

Convergence tests are shown in Table 6.1.3. Times are normalized to the fastest semi-Lagrangian solve. The ratio of successive changes in the solution (at a point) is reported.

BiCGSTAB with an ILUT preconditioner is used for the Solve routine (line 3 of Policy-Iteration) in this and all subsequent sections. In the specific case of the semi-Lagrangian scheme for the exchange rate problem, a simple tridiagonal solve can be used since the problem is a one-dimensional diffusion.

h	$u(t=0,x=0)$	Avg. policy its.	Avg. BiCGSTAB its.	Ratio	Norm. time
1	-0.60685256	3.13	0.74		1.32e+01
1/2	-0.61187228	2.88	0.90		6.99e+01
1/4	-0.61300925	2.58	0.93	4.42	3.98e+02
1/8	-0.61317577	2.49	0.94	6.83	2.77e+03
1/16	-0.61321292	2.48	0.95	4.48	2.09e+04
1/32	-0.61321903	2.46	0.95	6.08	1.61e+05

(a) Direct control

h	$u(t=0,x=0)$	Avg. policy its.	Avg. BiCGSTAB its.	Ratio	Norm. time
1	-0.60717652	3.19	0.71		1.38e+01
1/2	-0.61194960	2.88	0.76		6.96e+01
1/4	-0.61302973	2.55	0.91	4.42	3.95e+02
1/8	-0.61317966	2.48	1.28	7.20	2.76e+03
1/16	-0.61321390	2.48	1.33	4.38	2.09e+04
1/32	-0.61321928	2.46	0.99	6.36	1.61e+05

(b) Penalized

h	$u(t=0,x=0)$	Ratio	Norm. time
1	-0.69277804		1.00e+00
1/2	-0.64806716		6.49e+00
1/4	-0.62865965	2.30	4.90e+01
1/8	-0.62027822	2.32	3.86e+02
1/16	-0.61653511	2.24	3.17e+03
1/32	-0.61480123	2.16	2.64e+04
1/64	-0.61398311	2.12	2.17e+05

(c) Semi-Lagrangian

Table 6.1.3: Optimal combined control of the exchange rate: convergence tests

POLICY-ITERATION is terminated upon achieving a desired error tolerance:

$$\max_i \left\{ \frac{\left|v_i^k - v_i^{k-1}\right|}{\max\left(\left|v_i^k\right|, \text{scale}\right)} \right\} < \text{tol}.$$

The scale parameter ensures that unrealistic levels of accuracy are not imposed on the solution. We take $\text{tol} = 10^{-6}$ and $\text{scale} = 1$ for this and all future tests. The initial guess v^0 is taken to be the solution at the previous timestep, u^{n+1}.

Following [16], we take $\epsilon := D\Delta t$ and $\delta := 1/\epsilon$ with $D = 10^{-2}$.

For completeness, we mention that the obvious splitting with $\hat{\mu}(t,x) := 0$ and $\hat{f}(t,x) := -p(x)$ is used in the semi-Lagrangian scheme. The numerical examples of the sequel (§6.2 and §6.3) also use the obvious splittings.

The direct control and penalized schemes converge superlinearly. We speculate that this occurs since $x \mapsto u(t,x)$ is linear to the right of $x = \eta_0(t)$, and hence no error is made in approximating the term $D_x u$ and $D_x^2 u$ there. Assuming the solution u^{n+1} of the semi-Lagrangian scheme is linear to the right of $\eta_0(t^{n+1})$, error is introduced due to the approximation of $\eta_0(t^n)$ by $\eta_0(t^{n+1})$. This suggests that the direct control and penalized schemes may outperform the semi-Lagrangian scheme for problems with simple continuation regions and linear transaction costs.

Unsurprisingly, the direct control and penalized schemes are near-identical in performance and accuracy since the scaling and penalty factors are chosen identically (i.e. $\delta = 1/\epsilon$). We mention that the choice of $\delta = 1$ (i.e. no scaling) yields poor performance in the direct control setting (see [16] for an explanation).

Note that the average number of BiCGSTAB iterations per call to SOLVE can be less than one, suggesting that sometimes, no BiCGSTAB iterations are required on line 3 of POLICY-ITERATION. This occurs when the initial residual, $b(P^\ell) - A(P^\ell)v^{\ell-1}$, is small enough in magnitude (i.e. at the last policy iteration before convergence).

6.2 Optimal consumption and portfolio with both fixed and proportional transaction costs

The following is studied in [11]. Consider an investor that, at any point in time, has two investment opportunities: a stock and a bank account. Let $(S_t)_{t\geq 0}$ and $(B_t)_{t\geq 0}$ denote the amount of money invested in these two, respectively. The investor is able to

- consume continuously **(stochastic control)**;
- transfer money from the bank to the stock (or vice versa) subject to a transaction cost **(impulse control)**.

Denote by $(w_t)_{t\geq 0}$ the consumption rate with $0 \leq w_t \leq w_{\max}$. At any point in time, the investor can move money to ($z > 0$) or from ($z < 0$) the stock incurring a transaction cost of $\lambda|z| + C$ where $C > 0$ and $0 \leq \lambda < 1$. This is captured by

$$\begin{aligned}
dS_t &= \mu S_t dt + \xi S_t d\mathfrak{W}_t && \text{if } \tau_j < t < \tau_{j+1} && \text{(stochastic control);}\\
dB_t &= (rB_t - w_t)\, dt && \text{if } \tau_j < t < \tau_{j+1} && \text{(stochastic control);}\\
S_{\tau_{j+1}} &= S_{\tau_{j+1}-} + z_{\tau_{j+1}} && && \text{(impulse control);}\\
B_{\tau_{j+1}} &= B_{\tau_{j+1}-} - z_{\tau_{j+1}} - \lambda\left|z_{\tau_{j+1}}\right| - C && && \text{(impulse control).}
\end{aligned}$$

Parameter		Value	
Discount factor	ρ	10%	per annum
Interest rate	r	7%	per annum
Drift	μ	11%	per annum
Volatility	ξ	30%	per annum
Expiry	T	40	years
Relative risk aversion	$1-\gamma$	0.7	
Scaled transaction cost	λ	0.1	
Fixed transaction cost	C	0.05	
Maximum withdrawal rate	$w_{\max}$	100	
Initial stock value	s_0	\$45.20	
Initial bank account value	b_0	\$45.20	

Table 6.2.1: Optimal consumption: parameters from [11]

A combined control $\theta \coloneqq (w, \tau_1, \tau_2, \ldots, z_{\tau_1}, z_{\tau_2}, \ldots)$ is admissible if at all times, the stock holdings and bank account are nonnegative. Let Θ denote the set of all admissible controls.

The investor's maximal expected utility at time t with amount $S_t = s$ in the stock and $B_t = b$ in the bank account is given by

$$u(t,s,b) \coloneqq e^{\rho t} \sup_{\theta\in\Theta} \mathbb{E}^{(t,s,b)} \left[\int_t^T e^{-\rho t'} \frac{w_{t'}^\gamma}{\gamma} dt' + e^{-\rho T} \frac{\max\left(B_T + (1-\lambda) S_T - C, 0\right)^\gamma}{\gamma} \right]$$

where $0 \le 1-\gamma < 1$ is the investor's relative risk-aversion and $\rho \ge 0$ is the rate of time preference. The utility received at the expiry corresponds to liquidating the asset and consuming everything instantaneously.

The associated HJBQVI on $\Omega \coloneqq (0,\infty)^2$ and $\Lambda \coloneqq \emptyset$ is given by (1.1) with $g(T,x) \coloneqq \max(b + (1-\lambda)s - C, 0)^\gamma/\gamma$ and

$$W \coloneqq [0, w_{\max}]; \qquad Z(t,x) \coloneqq \left[-s, \frac{b-C}{1+\lambda}\right];$$

$$Ł^w \coloneqq \frac{1}{2}\xi^2 s^2 \frac{\partial^2}{\partial s^2} + \mu s \frac{\partial}{\partial s} + \begin{cases} (rb - w)\,\frac{\partial}{\partial b} & \text{if } b > 0; \\ 0 & \text{otherwise;} \end{cases} \qquad \Gamma(t,x,z) \coloneqq (z, -z - \lambda|z| - C);$$

$$f^w \coloneqq \begin{cases} w^\gamma/\gamma & \text{if } b > 0; \\ 0 & \text{otherwise;} \end{cases} \qquad K(t,x,z) \coloneqq 0.$$

In the above, expressions such as $s \cdot \partial/\partial s$ are to be interpreted as identically zero when $s = 0$. The convention $[q_1, q_2] = \emptyset$ if $q_1 > q_2$ is used.

6.2.1 Convergence of the direct control scheme

As in §6.1.1, the domain $[0,T] \times [0,\infty)^2$ and $Z(t,x)$ are truncated so that the state after an impulse $x_i + \Gamma(t, x_i, z_i)$ remains in the truncated domain. We use the notation $x_i = (s_i, b_i)$. The direct control problem is given by (4.3) subject to (5.1) and (5.2).

Suppose there exists a grid node x_{i^1} and $P \coloneqq (w, z, \psi)$ such that $\psi_{i^1} = 1$ and that there exists no path in $B(z)$ from i^1 to some j with $\psi_j = 0$. Since $C > 0$, there exists a path $i^1 \to i^2 \to \cdots$ of infinite length such that $s_{i^1} + b_{i^1} > s_{i^2} + b_{i^2} > \cdots$ and $\psi_{i^q} = 1$ for all q. Due to the finitude of the grid, $x_{i^q} = x_{i^\ell}$ (and hence $s_{i^q} + b_{i^q} = s_{i^\ell} + b_{i^\ell}$) for some $q < \ell$, a contradiction. It follows that no such x_{i^1} exists: (H2) is satisfied.

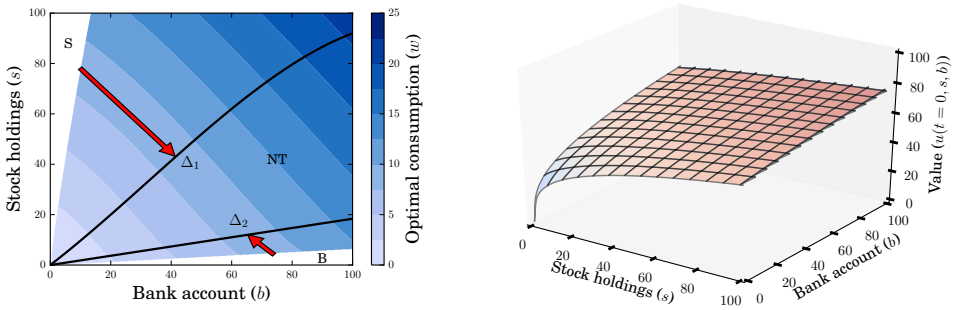

(a) Optimal control (b) Value

Figure 6.2.1: Optimal consumption at initial time (compare with [11, Figures 1 and 2])

h	s nodes	b nodes	w nodes	z nodes	Timesteps
1	20	20	15	15	32
1/2	40	40	30	30	64
⋮	⋮	⋮	⋮	⋮	⋮

Table 6.2.2: Optimal consumption: numerical grid

6.2.2 Optimal control

As in [11], three regions are observed in an optimal control: the *buy* (B), *sell* (S), and *continuation/no transaction* (NT) regions. In the B and S regions, the controller intervenes by jumping back to the closest of the two lines marked Δ_1 and Δ_2. In NT, the controller consumes continuously.

6.2.3 Convergence tests

Convergence tests are shown in Table 6.2.3. We mention that artificial Neumann boundary conditions $\partial^q u/\partial s^q = 0$ and $\partial u/\partial b = 0$ are used at the truncated boundaries $s = s_{\max}$ and $b = b_{\max}$. The results for the direct control and penalized schemes are near-identical, though the former requires significantly more policy iterations per timestep. The rate of convergence for the semi-Lagrangian scheme becomes sublinear for higher levels of refinement.

6.3 Guaranteed minimum withdrawal benefit (GMWB) in variable annuities

Guaranteed minimum withdrawal benefits (GMWB) in variable annuities provide investors with the tax-deferred nature of variable annuities along with a guaranteed minimum payment. GMWB pricing has been previously considered as a singular control problem in [21, 14] and as an impulse control problem in [12]. Optimal controls for GMWBs with annual withdrawals is considered in [1].

A GMWB is composed of investment and guarantee accounts, $(S_t)_{t\geq 0}$ and $(A_t)_{t\geq 0}$, respectively. It is bootstrapped via a lump sum payment s_0 to an insurer, placed in the

h	$u(t=0,s_0,b_0)$	Avg. policy its.	Avg. BiCGSTAB its.	Ratio	Norm. time
1	56.062123	7.63	1.28		1.63e+01
1/2	58.739224	8.80	1.90		2.93e+02
1/4	59.420125	10.4	2.28	3.93	5.66e+03
1/8	59.658413	11.8	3.28	2.86	1.03e+05
1/16	59.754780	13.3	4.45	2.47	1.85e+06
1/32	59.797206	14.2	6.54	2.27	3.05e+07

(a) Direct control

h	$u(t=0,s_0,b_0)$	Avg. policy its.	Avg. BiCGSTAB its.	Ratio	Norm. time
1	56.058496	4.09	1.43		1.03e+01
1/2	58.739041	3.95	1.77		1.47e+02
1/4	59.420075	3.40	2.35	3.94	1.99e+03
1/8	59.658399	3.04	3.69	2.86	2.69e+04
1/16	59.754778	2.80	5.50	2.47	4.02e+05
1/32	59.797215	2.58	5.98	2.27	5.86e+06

(b) Penalized

h	$u(t=0,s_0,b_0)$	Avg. BiCGSTAB its.	Ratio	Norm. time
1	55.621632	1.00		1.00e+00
1/2	58.782064	2.00		1.55e+01
1/4	59.404576	3.00	5.08	2.60e+02
1/8	59.569370	4.00	3.78	4.05e+03
1/16	59.651186	6.00	2.01	6.68e+04
1/32	59.705315	8.00	1.51	1.11e+06
1/64	59.748325	10.8	1.26	1.85e+07

(c) Semi-Lagrangian

Table 6.2.3: Optimal consumption: convergence tests

(risky) investment account (i.e. $S_0 = s_0$). A GMWB promises to pay back at least the lump sum s_0, assuming that the holder of the contract does not withdraw above a certain rate. This is captured by setting $A_0 = s_0$ and reducing both investment and guarantee accounts on a dollar-for-dollar basis upon withdrawals. The holder can continue to withdraw as long as the guarantee account remains positive. In particular, at any point in time until the expiry of the contract T, the holder may:

- withdraw continuously at a rate of $G \geq 0$ per annum regardless of the performance of the investment **(stochastic control)**;
- withdraw a finite amount z instantaneously reduced by the excess withdrawal rate $0 \leq \kappa \leq 1$ **(impulse control)**.

The holder gets the larger of the investment account and a full withdrawal at expiry.

The guarantee account can be withdrawn from continuously or instantaneously:

$$\begin{aligned} dA_t &= -w_t dt && \text{if } \tau_j < t < \tau_{j+1} && \text{(stochastic control);} \\ A_{\tau_{j+1}} &= A_{\tau_{j+1}-} - z_{\tau_{j+1}} && && \text{(impulse control).} \end{aligned}$$

Let $\rho \geq 0$ denote the risk-free rate. Consider an index $(Y_t)_{t \geq 0}$ following

$$dY_t = \rho Y_t dt + \xi Y_t d\mathfrak{W}_t$$

under the risk-neutral measure. The investment account tracks the index and is adjusted by withdrawals from the guarantee account:

$$\begin{aligned} dS_t &= \left((\rho - \eta)\, S_t - w_t \mathbb{1}_{\{S_t > 0\}}\right) dt + \xi S_t d\mathfrak{W}_t && \text{if } \tau_j < t < \tau_{j+1}; \\ S_{\tau_{j+1}} &= \max\left(S_{\tau_{j+1}-} - z_{\tau_{j+1}}, 0\right). \end{aligned}$$

$0 \leq \eta \leq \rho$ is the proportional rate deducted from the investment account and serves as a premium for the guarantee. A combined control $\theta \coloneqq (w, \tau_1, \tau_2, \ldots, z_1, z_2, \ldots)$ is admissible if at all times, the guarantee account is nonnegative. Let Θ denote the set of all admissible controls.

The insurer's worst-case cost of hedging (discussed in [2]) a GMWB at time t with amount $S_t = s$ in the risky account and amount $A_t = a$ is

$$u(t,s,a) \coloneqq e^{\rho t} \sup_{\theta \in \Theta} \mathbb{E}^{(t,s,a)} \Bigg[\int_t^T e^{-\rho t'} w_{t'} dt' + e^{-\rho T} \max\left(S_T, (1-\kappa)\, A_T - C\right) + \sum_{\tau_j \leq T} e^{-\rho \tau_j} \left((1-\kappa)\, z_{\tau_j} - C\right) \Bigg]$$

where $C > 0$ is a fixed transaction cost. The terminal payoff corresponds to the maximum of the investment account or withdrawing the entirety of the guarantee account at the excess withdrawal rate.

Let $x \coloneqq (s, a)$ and $\zeta \coloneqq \rho - \eta$. The associated HJBQVI on $\Omega \coloneqq (0, \infty)^2$ and $\Lambda \coloneqq \emptyset$ is given by (1.1) with $g(T, x) \coloneqq \max(s, (1-\kappa)a - C)$ and

Parameter		Value	
Risk-free rate	ρ	5%	per annum
Premium	η	0%	per annum
Volatility	ξ	30%	per annum
Expiry	T	10	years
Withdrawal rate	G	\$10	per annum
Excess withdrawal rate	κ	10%	
Fixed transaction cost	C	$1/10^6$	
Initial lump sum payment	s_0	\$100	

Table 6.3.1: GMWB: parameters from [12]

$$W \coloneqq [0, G]\,; \qquad Z(t,x) \coloneqq [0, a]\,;$$

$$Ł^w \coloneqq \frac{\xi^2 s^2}{2}\frac{\partial^2}{\partial s^2} + \zeta s \frac{\partial}{\partial s} + \begin{cases} -w\left(\frac{\partial}{\partial a} + \frac{\partial}{\partial s}\right) & \text{if } s, a > 0; \\ -w\frac{\partial}{\partial a} & \text{if } a > 0; \\ 0 & \text{otherwise;} \end{cases} \qquad \Gamma(t,x,z) \coloneqq -\left(\min\left(z, s\right), z\right);$$

$$f^w \coloneqq \begin{cases} w & \text{if } a > 0; \\ 0 & \text{otherwise;} \end{cases} \qquad K(t,x,z) \coloneqq (1-\kappa)\, z - C.$$

6.3.1 Convergence of the direct control scheme

We use the notation $x_i = (s_i, a_i)$ and assume the origin $(0,0)$ is part of the numerical grid. The direct control problem is given by (4.3) subject to (5.1) and (5.2).

Suppose (H4) is not satisfied so that for some solution v, there exists i such that $v_i = [\mathbb{B}v]_i = [\mathbb{B}^2 v]_i = \cdots$. Since $C > 0$, it follows that $v_i = -\infty$, a contradiction. Hence, (H4) holds.

We perform policy iteration on a modified problem with control set $\mathcal{P}'$ consisting of all controls $P \coloneqq (w, z, \psi)$ in $\mathcal{P}$ satisfying

$$\psi_i = 0 \text{ whenever } a_i = 0 \text{ and } z_i \neq 0 \text{ whenever } a_i \neq 0.$$

As in Example 4.5, (H2)′ follows from the unidirectionality of z_i. (4.7) is established by noting that $z_i = 0$ incurs an infinite cost (and is therefore suboptimal). Convergence then follows from an application of Theorem 4.10

Remark 6.2. The condition $z_i \neq 0$ appeals to intuition: the holder should never pay $C > 0$ for a withdrawal of zero dollars.

6.3.2 Optimal control

Figure 6.3.1 shows an optimal control for a GMWB, corresponding to a worst-case cost of hedging from the perspective of the insurer (optimality from the holder's perspective, who may have to take into consideration consumption, taxation, etc., is explored in [2]). We refer to [12] for an explanation of the three distinct withdrawal regions.

6.3.3 Convergence tests

Convergence tests are shown in Table 6.3.3. Since $w \mapsto Ł(t,x,w)$ is linear, we take $W_h = \{0, G\}$ independent of h. An asymptotic boundary condition is used at the truncated

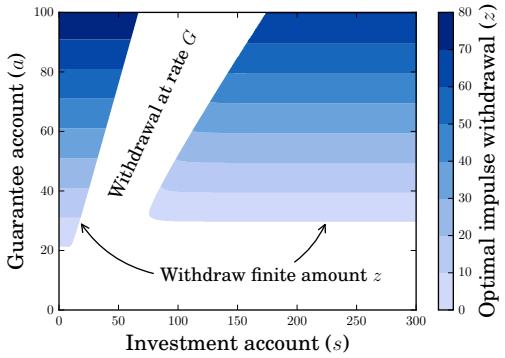

Figure 6.3.1: GMWB: optimal control at initial time with $\eta = 0.03126$ from [12]

h	w nodes	a nodes	z nodes	Timesteps
1	64	50	2	32
1/2	128	100	4	64
⋮	⋮	⋮	⋮	⋮

Table 6.3.2: GMWB: numerical grid

boundary $s = s_{\max}$ (no boundary condition is needed at $a = a_{\max}$ since the characteristics are outgoing in the a direction). For details, see [12]. The direct control and penalized scheme produce near-identical results and exhibit similar execution times.

7 Concluding remarks

This work establishes the well-posedness of (1.3) and gives sufficient conditions for convergence of the corresponding policy iteration. (1.3) has applications to the numerical solutions of HJBQVIs (§5–6) and infinite-horizon MDPs with vanishing discount factor (Corollary 4.4).

A semi-Lagrangian scheme for the HJBQVI (1.1) is both easy to implement and requires only one linear solve per timestep. However, it cannot be used if the diffusion or jump arrival rate of the underlying stochastic process are control-dependent.

The direct control and penalized schemes do not suffer these limitations. Numerical evidence suggests that both schemes perform similarly. However, policy iteration applied to the direct control scheme can fail (Example 4.9) unless additional care is taken to remove certain suboptimal controls. The removal of these controls is ad hoc (i.e. problem dependent). Therefore, we recommend discretizing the problem with a penalized scheme, applying policy iteration to solve the resulting nonlinear equations.

A General well-posedness of the Bellman problem (2.1)

By modifying policy iteration, it is possible to arrive at a version of Proposition 2.2 independent of (H1.ii). We can interpret this algorithm as taking into account the error from approximating the supremum in Policy-Iteration. The algorithm, closely related to [7,

h	$u(t=0,s_0,s_0)$	Avg. policy its.	Avg. BiCGSTAB its.	Ratio	Norm. time
1	107.68342	3.47	1.47		1.71e+01
1/2	107.70679	4.25	1.64		2.03e+02
1/4	107.71878	4.34	1.85	1.95	2.60e+03
1/8	107.72578	4.43	2.22	1.71	3.46e+04
1/16	107.72964	4.31	2.71	1.81	4.75e+05
1/32	107.73176	4.15	3.40	1.83	7.55e+06

(a) Direct control

h	$u(t=0,s_0,s_0)$	Avg. policy its.	Avg. BiCGSTAB its.	Ratio	Norm. time
1	107.68243	3.47	1.58		1.80e+01
1/2	107.70639	4.08	1.65		2.06e+02
1/4	107.71870	3.95	1.76	1.95	2.45e+03
1/8	107.72576	3.98	1.97	1.74	3.22e+04
1/16	107.72964	3.71	2.39	1.82	4.34e+05
1/32	107.73175	3.32	3.01	1.83	6.62e+06

(b) Penalized

h	$u(t=0,s_0,s_0)$	Avg. BiCGSTAB its.	Ratio	Norm. time
1	107.42351	1.00		1.00e+00
1/2	107.68443	1.00		1.02e+01
1/4	107.70841	1.00	10.9	1.39e+02
1/8	107.72257	1.00	1.70	2.03e+03
1/16	107.73015	1.00	1.87	3.31e+04
1/32	107.73224	1.98	3.62	6.10e+05
1/64	107.73337	2.90	1.85	1.13e+07

(c) Semi-Lagrangian

Table 6.3.3: GMWB: convergence tests

Algorithm Ho-4], is given below (subject to the convention that for x in $\mathbb{R}^M$ and c in $\mathbb{R}$, $x + c$ is the vector x with c added to each component):

SEQUENTIAL-POLICY-ITERATION$(\mathcal{P}, A(\cdot), b(\cdot), v^0)$

1 Pick a positive sequence $(\epsilon^\ell)_{\ell \geq 1}$ in $\mathbb{R}$ such that $\sum_{\ell \geq 1} \epsilon^\ell < \infty$
2 **for** $\ell = 1, 2, \ldots$
3 Pick P^ℓ such that $-A(P^\ell)v^{\ell-1} + b(P^\ell) + \epsilon^\ell \geq \sup_{P \in \mathcal{P}}\{-A(P)v^{\ell-1} + b(P)\}$
4 $v^\ell \coloneqq$ SOLVE$(A(P^\ell), b(P^\ell), v^{\ell-1})$

The following appears in [7]:

Lemma A.1. *A bounded sequence $(v^\ell)_{\ell \geq 0}$ in $\mathbb{R}$ converges if there exists a positive sequence $(\epsilon^\ell)_{\ell \geq 1}$ in $\mathbb{R}$ such that $\sum_{\ell \geq 1} \epsilon^\ell < \infty$ and $v^\ell - v^{\ell-1} \geq -\epsilon^\ell$ for $\ell \geq 1$.*

We require the following lemma, whose proof is trivial and thus omitted:

Lemma A.2. *Let X be a set, Y a normed linear space, $T\colon X \times Y \to \mathbb{R}$, and $Q\colon X \to \mathbb{R}$ with Q bounded above. Suppose that for each x in X, $T_x\colon Y \to \mathbb{R}$ defined by $T_x(y) \coloneqq T(x, y)$ is linear and that T_x has operator norm bounded uniformly with respect to x. The map $y \mapsto \sup_{x \in X}\{T(x, y) + Q(x)\}$ is uniformly continuous.*

Theorem A.3 (Convergence of sequential policy iteration)**.** *Suppose (H0), (H1.i), and that $A(P)$ is a monotone matrix for all P in $\mathcal{P}$. $(v^\ell)_{\ell \geq 1}$ defined by* SEQUENTIAL-POLICY-ITERATION *converges to the unique solution v of (2.1).*

Proof. First, note that

$$A(P^\ell)\left(v^\ell - v^{\ell-1}\right) = -A(P^\ell)v^{\ell-1} + b(P^\ell) \geq \sup_{P \in \mathcal{P}}\left\{-A(P)v^{\ell-1} + b(P)\right\} - \epsilon^\ell. \tag{A.1}$$

For $\ell > 1$,

$$\sup_{P \in \mathcal{P}}\left\{-A(P)v^{\ell-1} + b(P)\right\} \geq -A(P^{\ell-1})v^{\ell-1} + b(P^{\ell-1}) = 0.$$

Combining this with (A.1),

$$v^\ell - v^{\ell-1} \geq -A(P^\ell)^{-1}\left(\epsilon^\ell, \ldots, \epsilon^\ell\right)^\intercal \geq -C\epsilon^\ell \text{ for some } C \geq 0.$$

The last inequality follows from the boundedness of $P \mapsto A(P)^{-1}$ in (H0). By Lemma A.1, $v^\ell \to v$ for some v in $\mathbb{R}^M$. Taking limits in (A.1) and applying Lemma A.2,

$$0 = \lim_{\ell \to \infty}\left(\sup_{P \in \mathcal{P}}\left\{-A(P)v^{\ell-1} + b(P)\right\}\right) = \sup_{P \in \mathcal{P}}\left\{-A(P)v + b(P)\right\}.$$

Hence, v is a solution to (2.1). Uniqueness is proven similarly to Theorem 4.8. ∎

B Proof of Lemma 4.1

Proof. We write A_δ and b_δ to stress dependence on δ. Let v be a solution of (4.3) with $\delta = 1$. A pigeonhole principle argument allows us to pick a sequence $(P^\ell)_{\ell \geq 0} \coloneqq (w^\ell, z^\ell, \psi^\ell)_{\ell \geq 0}$ in $\mathcal{P}$ such that $\psi^\ell = \psi$ is constant and $-A_1(P^\ell)v + b_1(P^\ell) \to 0$. Multiplying both sides by $I - \Psi + \delta_0\Psi$ (where $\Psi \coloneqq \operatorname{diag}(\psi)$) yields $-A_{\delta_0}(P^\ell)v + b_{\delta_0}(P^\ell) \to 0$, and hence $\sup_{P \in \mathcal{P}}\{-A_{\delta_0}(P)v + b_{\delta_0}(P)\} \geq 0$. Supposing that this inequality is strict, it follows that for some P, $-A_{\delta_0}(P)v + b_{\delta_0}(P) > 0$. Multiplying both sides by $I - \Psi + \delta_0^{-1}\Psi$ yields $-A_1(P)v + b_1(P) > 0$, contradicting that v is a solution. The converse is handled similarly. ∎

References

[1] P. Azimzadeh and P. A. Forsyth. The existence of optimal bang-bang controls for GMxB contracts. *SIAM J. Financial Math.*, 6(1):117–139, 2015.

[2] P. Azimzadeh, P. A. Forsyth, and K. R. Vetzal. Hedging costs for variable annuities under regime-switching. In *Hidden Markov Models in Finance*, pages 133–166. Springer, 2014.

[3] J. Babbin, P. A. Forsyth, and G. Labahn. A comparison of iterated optimal stopping and local policy iteration for American options under regime switching. *J. Sci. Comput.*, 58(2):409–430, 2014.

[4] G. Barles and P. E. Souganidis. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptot. Anal.*, 4(3):271–283, 1991.

[5] E. Bayraktar and H. Xing. Pricing Asian options for jump diffusion. *Math. Finance*, 21(1):117–143, 2011.

[6] A. Bensoussan and J. L. Lions. *Impulse control and quasi-variational inequalities*. Gaunthier-Villars, Paris, 1984.

[7] O. Bokanowski, S. Maroso, and H. Zidani. Some convergence results for Howard's algorithm. *SIAM J. Numer. Anal.*, 47(4):3001–3026, 2009.

[8] J. H. Bramble and B. E. Hubbard. On a finite difference analogue of an elliptic boundary problem which is neither diagonally dominant nor of non-negative type. *J. Math. Phys.*, 43(2):117, 1964.

[9] A. Cadenillas and F. Zapatero. Optimal central bank intervention in the foreign exchange market. *J. Econom. Theory*, 87(1):218–242, 1999.

[10] J. P. Chancelier, M. Messaoud, and A. Sulem. A policy iteration algorithm for fixed point problems with nonexpansive operators. *Math. Methods Oper. Res.*, 65(2):239–259, 2007.

[11] J. P. Chancelier, B. Øksendal, and A. Sulem. Combined stochastic control and optimal stopping, and application to numerical approximation of combined stochastic and impulse control. *Proc. Steklov Inst. Math.*, 237(0):149–172, 2002.

[12] Z. Chen and P. A. Forsyth. A numerical scheme for the impulse control formulation for pricing variable annuities with a guaranteed minimum withdrawal benefit (GMWB). *Numer. Math.*, 109(4):535–569, 2008.

[13] R. Cont and E. Voltchkova. A finite difference scheme for option pricing in jump diffusion and exponential lévy models. *SIAM J. Numer. Anal.*, 43(4):1596–1626, 2005.

[14] M. Dai, Y. K. Kwok, and J. Zong. Guaranteed minimum withdrawal benefit in variable annuities. *Math. Finance*, 18(4):595–611, 2008.

[15] P. A. Forsyth and G. Labahn. Numerical methods for controlled Hamilton-Jacobi-Bellman PDEs in finance. *J. Comput. Finance*, 11(2):1, 2007.

[16] Y. Huang, P. A. Forsyth, and G. Labahn. Inexact arithmetic considerations for direct control and penalty methods: American options under jump diffusion. *Appl. Numer. Math.*, 72:33–51, 2013.

[17] K. Ishii. Viscosity solutions of nonlinear second order elliptic PDEs associated with impulse control problems. *Funkcial. Ekvac.*, 36(1):123–141, 1993.

[18] I. Kharroubi, J. Ma, H. Pham, J. Zhang, et al. Backward SDEs with constrained jumps and quasi-variational inequalities. *Ann. Appl. Probab.*, 38(2):794–840, 2010.

[19] H. J. Kushner and P. G. Dupuis. *Numerical methods for stochastic control problems in continuous time*. Springer-Verlag, New York, 1992.

[20] H. Le and C. Wang. A finite time horizon optimal stopping problem with regime switching. *SIAM J. Control Optim.*, 48(8):5193–5213, 2010.

[21] M. A. Milevsky and T. S. Salisbury. Financial valuation of guaranteed minimum withdrawal benefits. *Insurance Math. Econom.*, 38(1):21–38, 2006.

[22] G. Mundaca and B. Øksendal. Optimal stochastic intervention control with application to the exchange rate. *J. Math. Econom.*, 29(2):225–243, 1998.

[23] A. M. Oberman. Convergent difference schemes for degenerate elliptic and parabolic equations: Hamilton-Jacobi equations and free boundary problems. *SIAM J. Numer. Anal.*, 44(2):879–895, 2006.

[24] B. Øksendal and A. Sulem. *Applied stochastic control of jump diffusions*, volume 498. Springer, 2005.

[25] R. J. Plemmons. M-matrix characterizations.I–Nonsingular M-matrices. *Linear Algebra Appl.*, 18(2):175–188, 1977.

[26] R. C. Seydel. Existence and uniqueness of viscosity solutions for QVI associated with impulse control of jump-diffusions. *Stochastic Process. Appl.*, 119(10):3719–3748, 2009.

[27] P. N. Shivakumar and K. H. Chew. A sufficient condition for nonvanishing of determinants. *Proc. Amer. Math. Soc.*, pages 63–66, 1974.

[28] J. H. Witte and C. Reisinger. Penalty methods for the solution of discrete HJB equations-continuous control and obstacle problems. *SIAM J. Numer. Anal.*, 50(2):595–625, 2012.